psyc3010 lecture 5

consolidation of between-subjects anova
power

last week: complex anova
next week: Blocking and ANCOVA

last week → this week

- last week we looked at how a 2-way anova can be extended to a 3-way anova, where three factors (IVs) are included in the model
- this week we discuss the assignment
- briefly recap on 3-way anova
  - Including simple interactions, simple simple effects, and simple simple comparisons
  - A quick note on unequal n
- Then on to power
  - What it is, graphically & statistically
  - effect size measures, d and Φ'(phi prime)
  - non-centrality parameters δ (delta) and Φ (phi)
  - What influences power

last week's experiment with some ***different*** data…

Males

Personality	Reinforcement Rew	None	Pun
Impulsive	350	425	450
	345	430	465
	330	445	465
Total	***1025***	***1300***	***1380***
Mean	***341.67***	***433.33***	***460.00***
Anxious	485	450	470
	450	455	480
	495	445	465
Total	***1430***	***1350***	***1415***
Mean	***476.67***	***450.00***	***471.67***

Females

Personality	Reinforcement Rew	None	Pun
Impulsive	350	355	460
	320	350	495
	330	340	485
Total	***1000***	***1045***	***1440***
Mean	***333.33***	***348.33***	***480.00***
Anxious	320	360	480
	340	360	470
	350	370	450
Total	***1010***	***1090***	***1400***
Mean	***336.67***	***363.33***	***466.67***

omnibus tests in a 3-way anova

- **main effects:**
  - differences between marginal means of one factor (averaging over levels of other factors)
- **two-way interactions (first-order):**
  - examines whether the effect of one factor is the same at every level of a second factor (averaging over a third factor)
- ***three-way interaction* (second-order)*:***
  - examines whether the two-way interaction between two factors is the same at every level of the third factor
  - *or* whether the cell means differ more than you would expect given the main effects and the two-way interactions

main summary table (from SPPS)

Tests of Between-Subjects Effects

Dependent Variable: RT

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
REINFORC	60237.500	2	30118.750	169.418	.000
PERSONAL	7802.778	1	7802.778	43.891	.000
GENDER	24544.444	1	24544.444	138.062	.000
REINFORC * PERSONAL	7751.389	2	3875.694	21.801	.000
REINFORC * GENDER	16426.389	2	8213.194	46.199	.000
PERSONAL * GENDER	6944.444	1	6944.444	39.062	.000
REINFORC * PERSONAL * GENDER	6601.389	2	3300.694	18.566	.000
Error	4266.667	24	177.778		
Total	134575.000	35			

steps for following-up a 3-way interaction

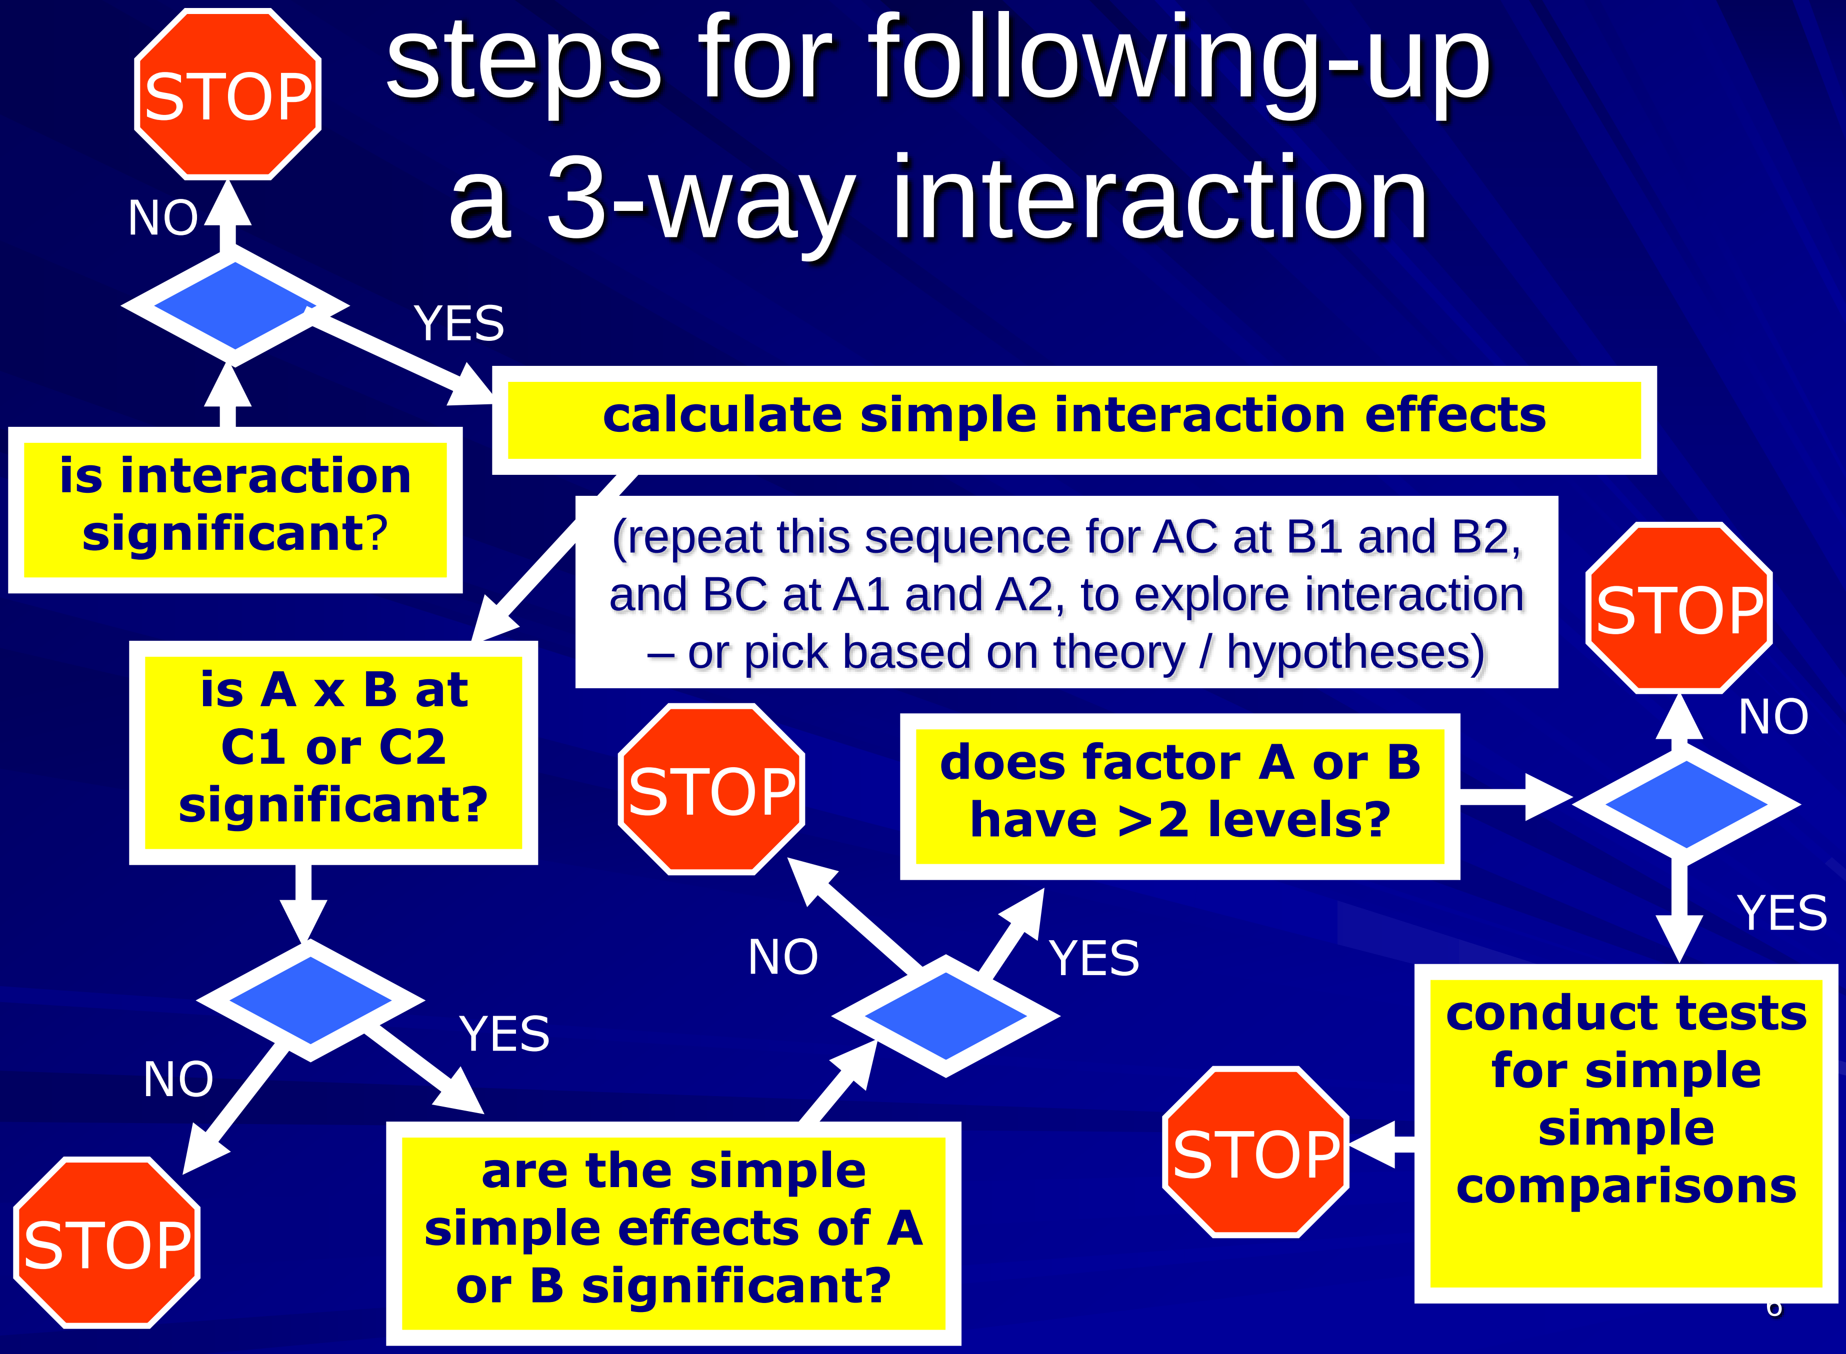

following up the 3-way interaction

- **simple interaction effects**
  – tests of 2-way interactions at each level of the third factor
  – e.g., personality x reinforcement at each level of gender

- **simple simple effects:**
  – tests of the effect of a factor at each level of second factor within levels of a third factor
  – e.g., effects of reinforcement at each level of personality (for males)

- **simple simple comparisons:**
  – compares effect of different levels of a factor at a particular level of a second factor within levels of a third factor
  – e.g., effects of rew versus pun or neutral for impulsive males

summary table for follow-up tests

Source	*SS*	*df*	*MS*	*F*	*p*
Simple Interaction Effects					
PR at G1	13744.44	2	6872.22	38.656	0.000
PR at G2	608.33	2	304.17	1.711	0.202
Error	4266.67	24	177.78		
Simple Simple effects in the interaction of PR at G1					
R at P1 at G1	23116.67	2	11558.33	65.016	0.000
R at P2 at G1	538.89	2	269.44	1.516	0.240
Error	4266.67	24	177.78		
Simple Simple comparisons in simple effect of R at P1 at G1					
R vs N and P at P1 at G1	22047.90	1	22047.90	124.019	0.000
N vs P at P1 at G1	355.64	1	355.64	2.000	0.170
Error	4266.67	24	177.78		

summary table for follow-up tests

Source	*SS*	*df*	*MS*	*F*	*p*
Simple Interaction Effects					
PR at G1	13744.44	2	6872.22	38.656	0.000
PR at G2	608.33	2	304.17	1.711	0.202
Error	4266.67	24	177.78		

the personality x reinforcement interaction is…

- significant for males

- non-significant for females

we can follow up this significant simple interaction effect of Personality x Reinforcement for males using ***simple simple effects***

summary table for follow-up tests

the effect of reinforcement (for males) is

- significant for the impulsive group
- not-significant for the anxious group

Source	*SS*	*df*	*MS*	*F*	*p*
Simple Simple effects in the interaction of PR at G1					
R at P1 at G1	23116.67	2	11558.33	65.016	0.000
R at P2 at G1	538.89	2	269.44	1.516	0.240
Error	4266.67	24	177.78		

we can follow up the significant simple simple effect of reinforcement at impulsivity (for males) using ***simple simple comparisons***

summary table for follow-up tests

simple simple comparisons…

- calculated using the t-test formula (c.f. lecture 3)

- can also be calculated using a F-test

(last week's tutorial exercise showed how to do this)

Source	*SS*	*df*	*MS*	*F*	*p*
Simple Simple comparisons in simple effect of R at P1 at G1					
R vs N and P at P1 at G1	22047.90	1	22047.90	124.019	0.000
N vs P at P1 at G1	355.64	1	355.64	2.000	0.170
Error	4266.67	24	177.78		

summary table for follow-up tests

for impulsive males,

1) performance under ***rewarding*** feedback was significantly different from that under ***neutral*** or ***punishing*** feedback

2) performance under ***punishing*** feedback was not significantly different to performance under ***neutral*** feedback

Simple Simple comparisons in simple effect of R at P1 at G1

Source	***SS***	***df***	***MS***	***F***	***p***
R vs N and P at P1 at G1	22047.90	1	22047.9	124.019	0.000
N vs P at P1 at G1	355.64	1	355.644	2.000	0.170
Error	4266.67	24	177.78		

reporting

1. I haven't put effect sizes. These would be required for all omnibus tests. Effect sizes are usually included for simple effects as well, and even for simple comparisons.

A 2 x 2 x 3 factorial ANOVA was performed, and a significant three-way interaction was observed among personality, gender, and reinforcement $F(2,24) = 169.42$, $p<.001$. The simple personality x reinforcement was significant for men, $F(2,24) = 38.66$, $p<.001$, but not for women, $F(2,24) = 1.71$, $p = .202$. Simple simple effects of reinforcement were significant for impulsive males, $F(2,24) = 65.02$, $p<.001$, but not for anxious males, $F(2,24) = 1.52$, $p = .240$. Planned contrasts showed that impulsive males performed better with reward ($M = 341.67$) than with punishment or no reinforcement ($M = 441.67$), $F(1,24) = 124.02$, $p<.001$, while the performance of impulsive males receiving no reinforcement ($M = 433.33$) was not significantly different to those receiving punishment ($M = 460.00$), $F(2,24) = 2.00$, $p = .170$.

reporting

2. NB we are reporting simple simple effects of reinforcement now – assumes this = theoretical focus. Last week we hypothesized different personality effects depending on reinforcement, but this week we are focusing on different reinforcement effects depending on personality.

Results of a 2 x 2 x 3 factorial ANOVA indicated a significant three-way interaction among personality, gender, and reinforcement $F(2,24) = 169.42$, $p<.001$. The simple personality x reinforcement interaction was significant for males, $F(2,24) = 38.66$, $p<.001$, but not for females, $F(2,24) = 1.71$, $p = .202$. There was a significant simple simple effect of reinforcement for impulsive males, $F(2,24) = 65.02$, $p<.001$, but not for anxious males, $F(2,24) = 1.52$, $p = .240$. Planned contrasts showed that impulsive males performed better with reward ($M = 341.67$) than with punishment or no reinforcement ($M = 441.67$), $F(1,24) = 124.02$, $p<.001$, while the performance of impulsive males receiving no reinforcement ($M = 433.33$) was not significantly different to those receiving punishment ($M = 460.00$), $F(2,24) = 2.00$, $p = .170$.

reporting

3. NB for simplicity we reported for women that the simple interaction of P x R was ns, but unless interested ONLY in interaction we would also ‘followup’ ns simple interaction by looking at simple effect of Reinforcement for women (collapsing across Personality), if hypotheses suggest R = the theoretically key IV. We want to know if there are effects of reinforcement for women that don’t change depending on personality.

Results of a 2 x 2 x 3 factorial ANOVA indicated a significant three-way interaction among personality, gender, and reinforcement $F(2,24) = 169.42$, $p<.001$. The simple personality x reinforcement interaction was significant for males, $F(2,24) = 38.66$, $p<.001$, but not for females, $F(2,24) = 1.71$, $p = .202$. [Insert before simple simple effects of R for men:] For women, across personality type the simple effect of reinforcement was / was not significant, F(2,24)=X, p=, eta2=…. [If yes, would also follow-up with contrasts.] In addition, there was a significant simple simple effect of reinforcement for impulsive males, $F(2,24) = 65.02$, $p<.001$, but not for anxious males, $F(2,24) = 1.52$, $p = .240$. etc”

Reporting

- Where there is one key IV with two moderators of interest, you might also choose to skip directly from a three-way to the simple simple effects.
  - E.g. (made-up data), “The three-way interaction was significant, F=,p=,eta2=. Follow-up tests examined the simple simple effects of reinforcement within each level of personality and gender. Differences in reaction to reinforcement were observed for anxious men, F=,p=,eta2=, and impulsive women,F=,p=,eta2=. Specifically, [describe simple simple comparisons here]. For impulsive men and anxious women, no effects of reinforcement were found, Fs< ##, ps> ##, eta2s< ##.”
- In fact you might even go directly from the three-way to the simple comparisons (and put these in a table) to reduce your wordiness and the complexity of the text.

Table 1. Mean reaction time as a function of gender, personality, and reinforcement.

1. Most experimental journals would also want to see standard deviations.

Reinforcement:	Punishment	Neutral	Reinforcement
Women			
Impulsive	485.33_a	481.67_a	320.00_b
Anxious	487.10_a	485.33_a	323.00_b
Men			
Impulsive	484.25_a	478.76_a	316.67_b
Anxious	367.55_a	455.67_b	583.00_c

Note. Differing subscripts within the same row indicate significant differences among the means.

E.g. (made-up data): “The three-way interaction of P x R x G was significant, F=,p=,eta2=. Table 1 displays the means by gender, personality, and reinforcement to decompose the interaction. As seen in Table 1, follow-up analyses revealed that for women and for impulsive men, reinforcement led to faster responses (more learning) than either the control condition or punishment, which did not differ. For anxious men, by contrast, punishment produced the most learning, followed by the control condition, with reinforcement being the least successful condition (slowest response times).”

Table 1. Mean reaction time as a function of gender, personality, and reinforcement.

2. NB how this hides all stats (except the three-way) – excellent for non-stats minded reviewers

Reinforcement:	Punishment	Neutral	Reinforcement
Women			
Impulsive	485.33_a	481.67_a	320.00_b
Anxious	487.10_a	485.33_a	323.00_b
Men			
Impulsive	484.25_a	478.76_a	316.67_b
Anxious	367.55_a	455.67_b	583.00_c

Note. Differing subscripts within the same row indicate significant differences among the means.

E.g. (made-up data): "The three-way interaction of P x R x G was significant, F=,p=,eta2=. Table 1 displays the means by gender, personality, and reinforcement to decompose the interaction. As seen in Table 1, follow-up analyses revealed that for women and for impulsive men, reinforcement led to faster responses (more learning) than either the control condition or punishment, which did not differ. For impulsive men, by contrast, punishment produced the most learning, followed by the control condition, with reinforcement being the least successful condition (slowest response times)."

Reporting: The bottom line

- With higher-order designs, there’s a lot more flexibility about how you structure the analyses / write-up
- It’s best to be guided by your theory.
- This may require you to fully explore the simple interactions, simple simple effects, simple simple comparisons, etc. within the text…
- Or allow for some skipping.
- ➢ The more complex the design, the more likely that you need a Table (or Tables, and/or figures) to report the results.

What happens with Unequal n?

- Something really complex mathematically – so you're not required to learn this material (it will not be assessed on the exam)
- We provide a quick overview of the issues for the interested among you on the Blackboard web site in the practice materials section
- **Take home message:** If unequal n are too extreme (ratio of largest to smallest cell size > 3:1), results of analyses are unstable (less likely to replicate – can have more Type 1 *and* more Type 2 errors).
  - Always try for equal n in study designs where anticipate using ANOVA.
  - Avoid using natural group variables like gender or ethnicity in ANOVA if ratio of people in groups exceeds 3:1 for largest: smallest.

power - introduction

- See Howell sec. 11.13, 13.7
- significant differences are defined with reference to a **criterion**, or **threshold** (controlled/acceptable rate) for committing type-1 errors, typically .05 or .01
  - **the type-1 error** → finding a significant difference in the sample when it actually doesn't exist in the population
  - type-1 error rate denoted α
- in the development of hypothesis testing procedures, little attention has been paid to the type-2 error
  - **the type-2 error** → finding no significant difference in the sample when there is a difference in the population
  - type-2 error rate denoted β

power

shifts our focus to the type-2 error rate:

official definition: "the probability of correctly rejecting a false H_0" - mathematically works out to $1 - \beta$

useful definition: the degree to which we can detect treatment effects

(incl main effects, interactions, simple simple comparisons etc)

Reality vs Statistical Decisions

Statistical Decision:	Reality: H0	Reality: H1
Reject H0	"False alarm" α	"Hit" (grooviness!) 1 - β Power
Fail to reject H0	"Hit" (grooviness!) 1- α	"Miss" β

power

Common constraints :

–Cell size too small

- B/c sample difficult to recruit or too little time / money

–Small effects are often a focus of theoretical interest (especially in social / clinical / org – everything “real world”)

- DV is subject to multiple influences, so each IV has small impact
- “Error” or residual variance is large, because many IVs unmeasured in experiment / survey are influencing DV
- Interactions are of interest, and interactions draw on smaller cell sizes (and thus lower power) than tests of main effects [Cell means for interaction are based on n observations, while main effects are based on n x # of levels of other factors collapsed across]

power – kinds of questions

1) "I have done a study and I want to report the power of my significant effect"
   – need to calculate *observed power (post hoc power)*
2) "I have done a study and did not find a significant effect: how could I increase power to detect the difference?"

or

3) "I am designing a study and want to know how many subjects I need in order to detect my predicted effect"
   – need to calculate *predicted power (a priori power)*

in terms of hypothesis testing

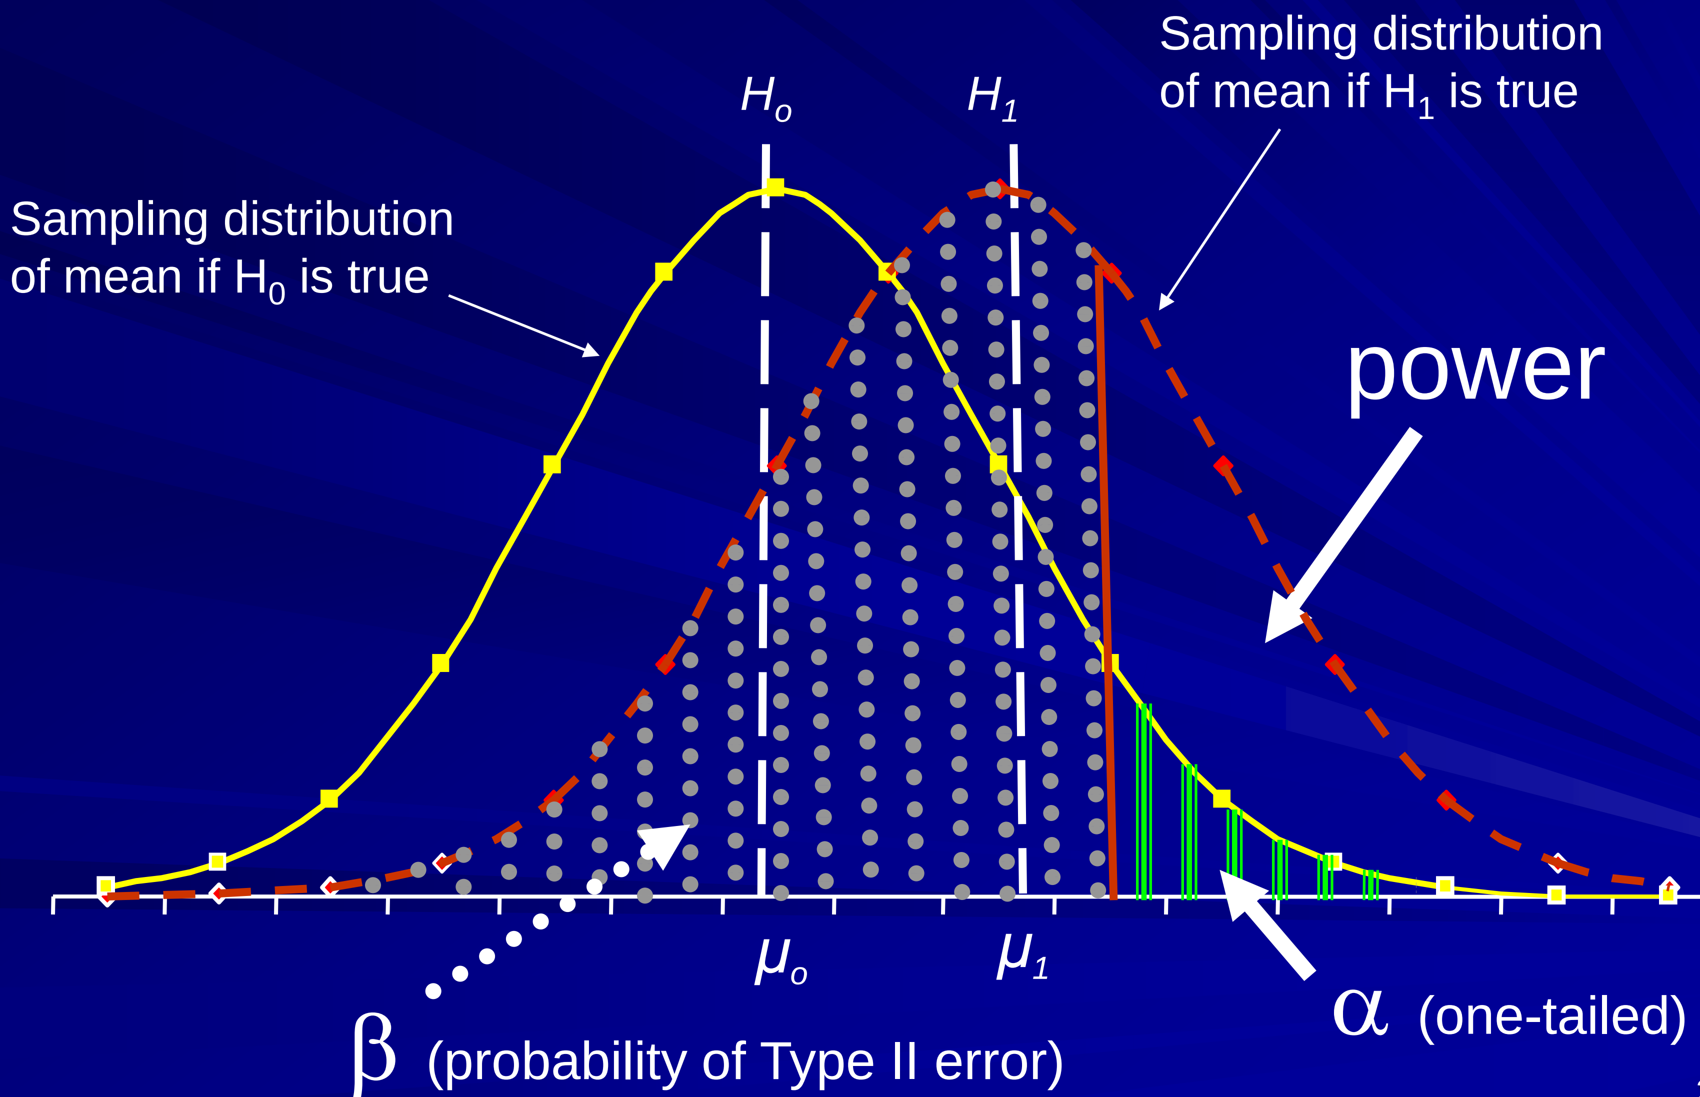

factors affecting power

- **power depends on a number of things**
  - significance level, α
    - relaxed α → more power
  - sample size, **N**
    - more N → more power
  - mean differences, $\mu_0 - \mu_1$
    - larger differences → more power
  - error variance – σ_e^2 **or** $\mathbf{MS_{error}}$
    - less error variance → more power

power as a function of α

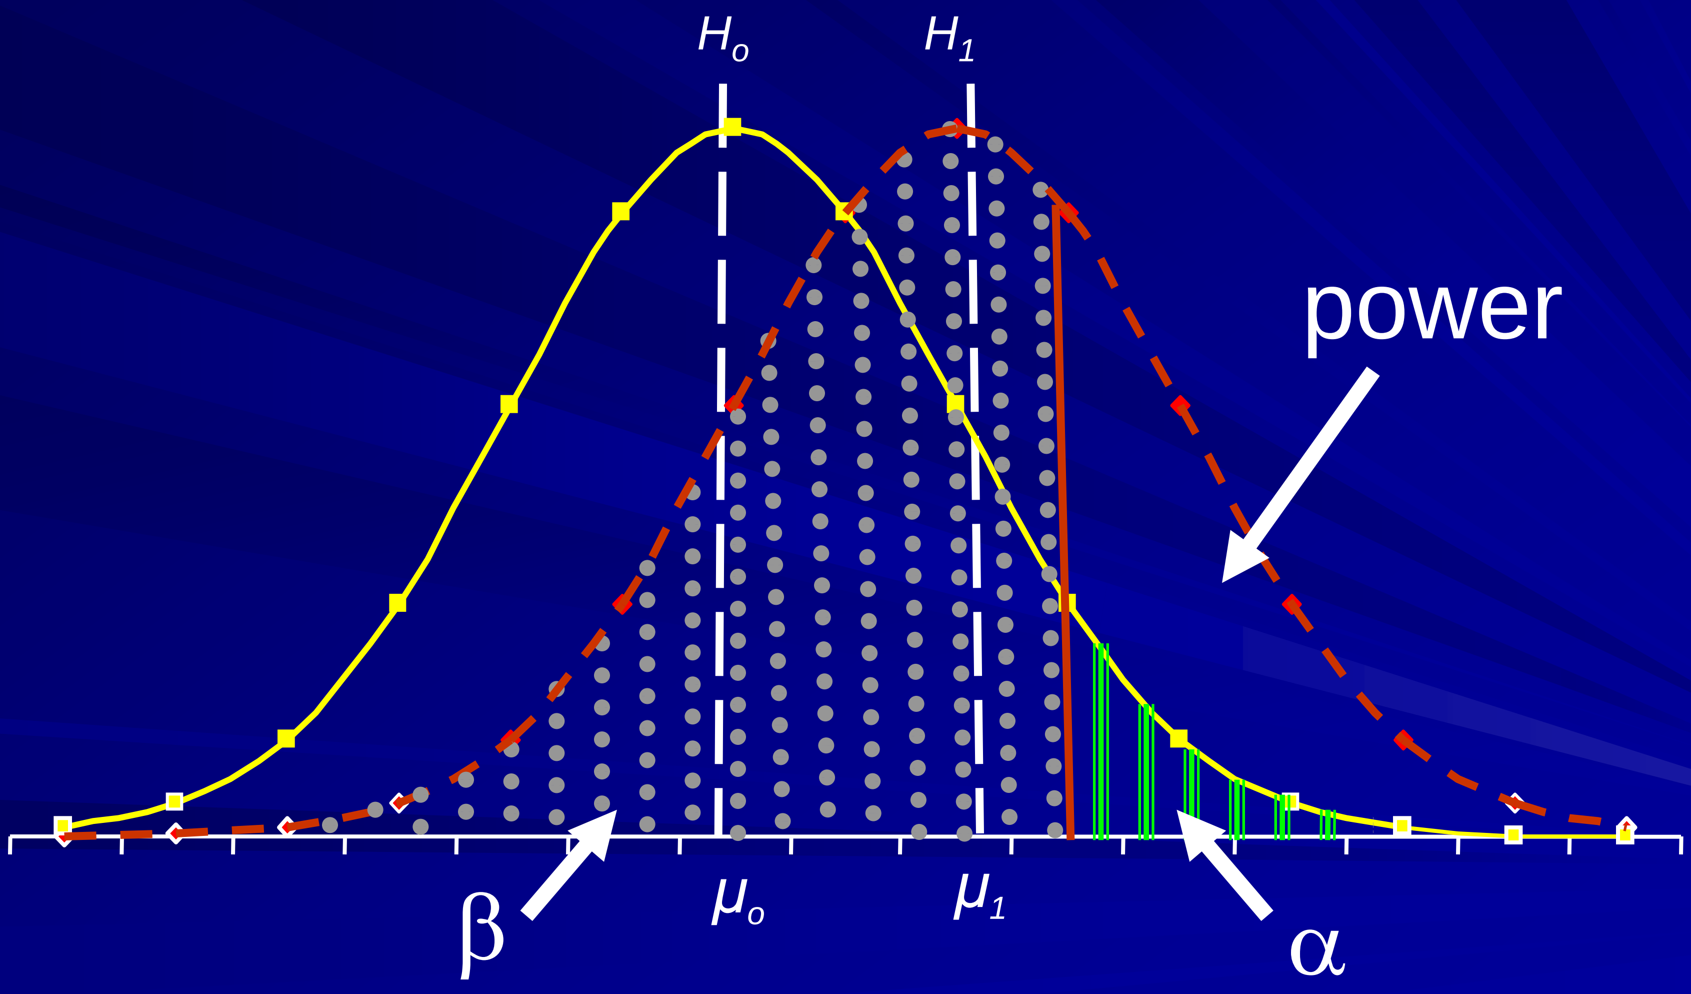

power as a function of α

occasionally people will set α=.10 if power is low

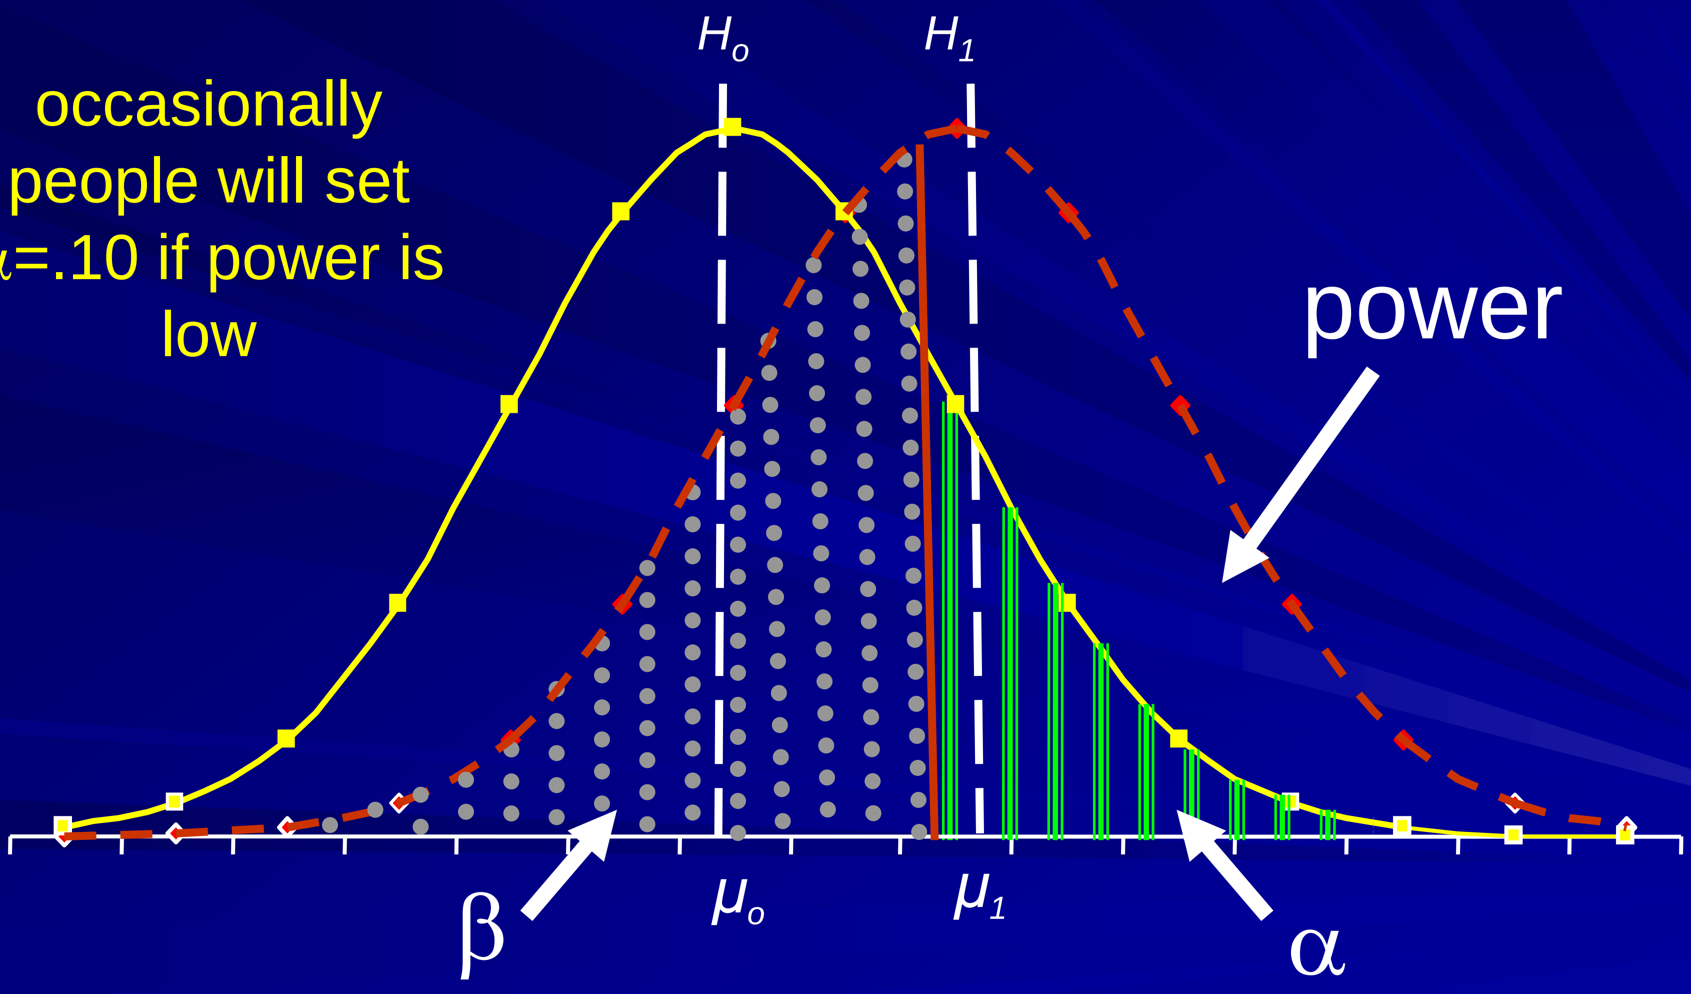

power as a function of $\mu_0 - \mu_1$

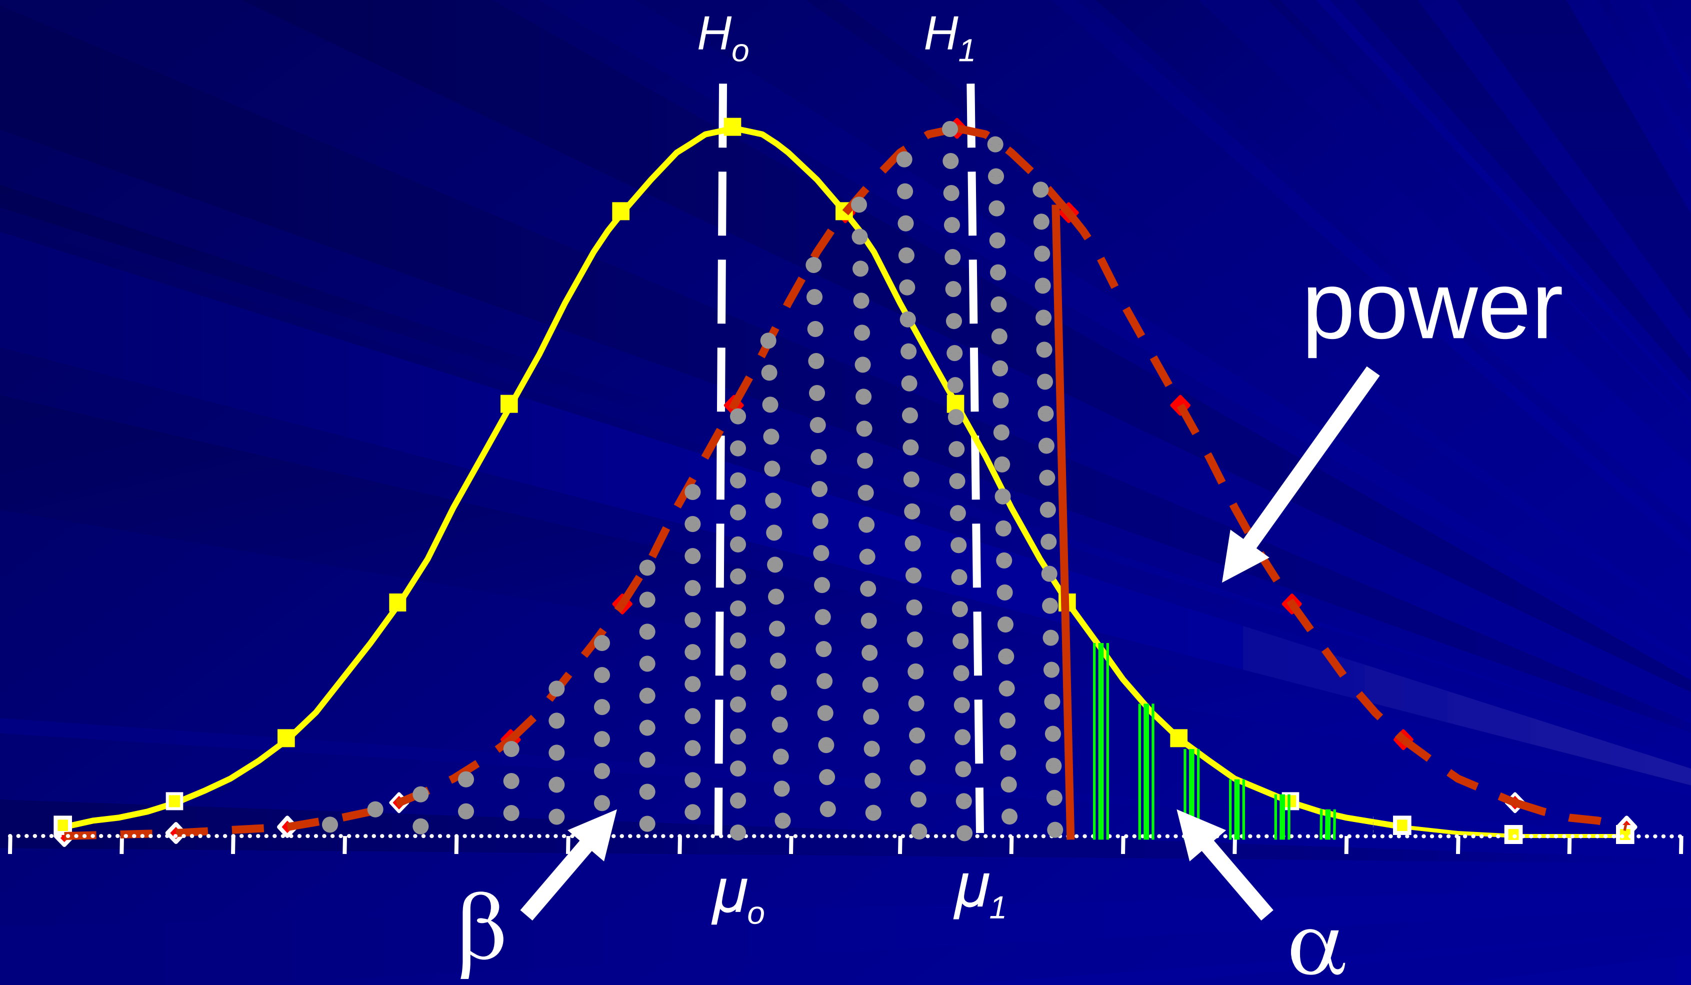

power as a function of $\mu_0 - \mu_1$

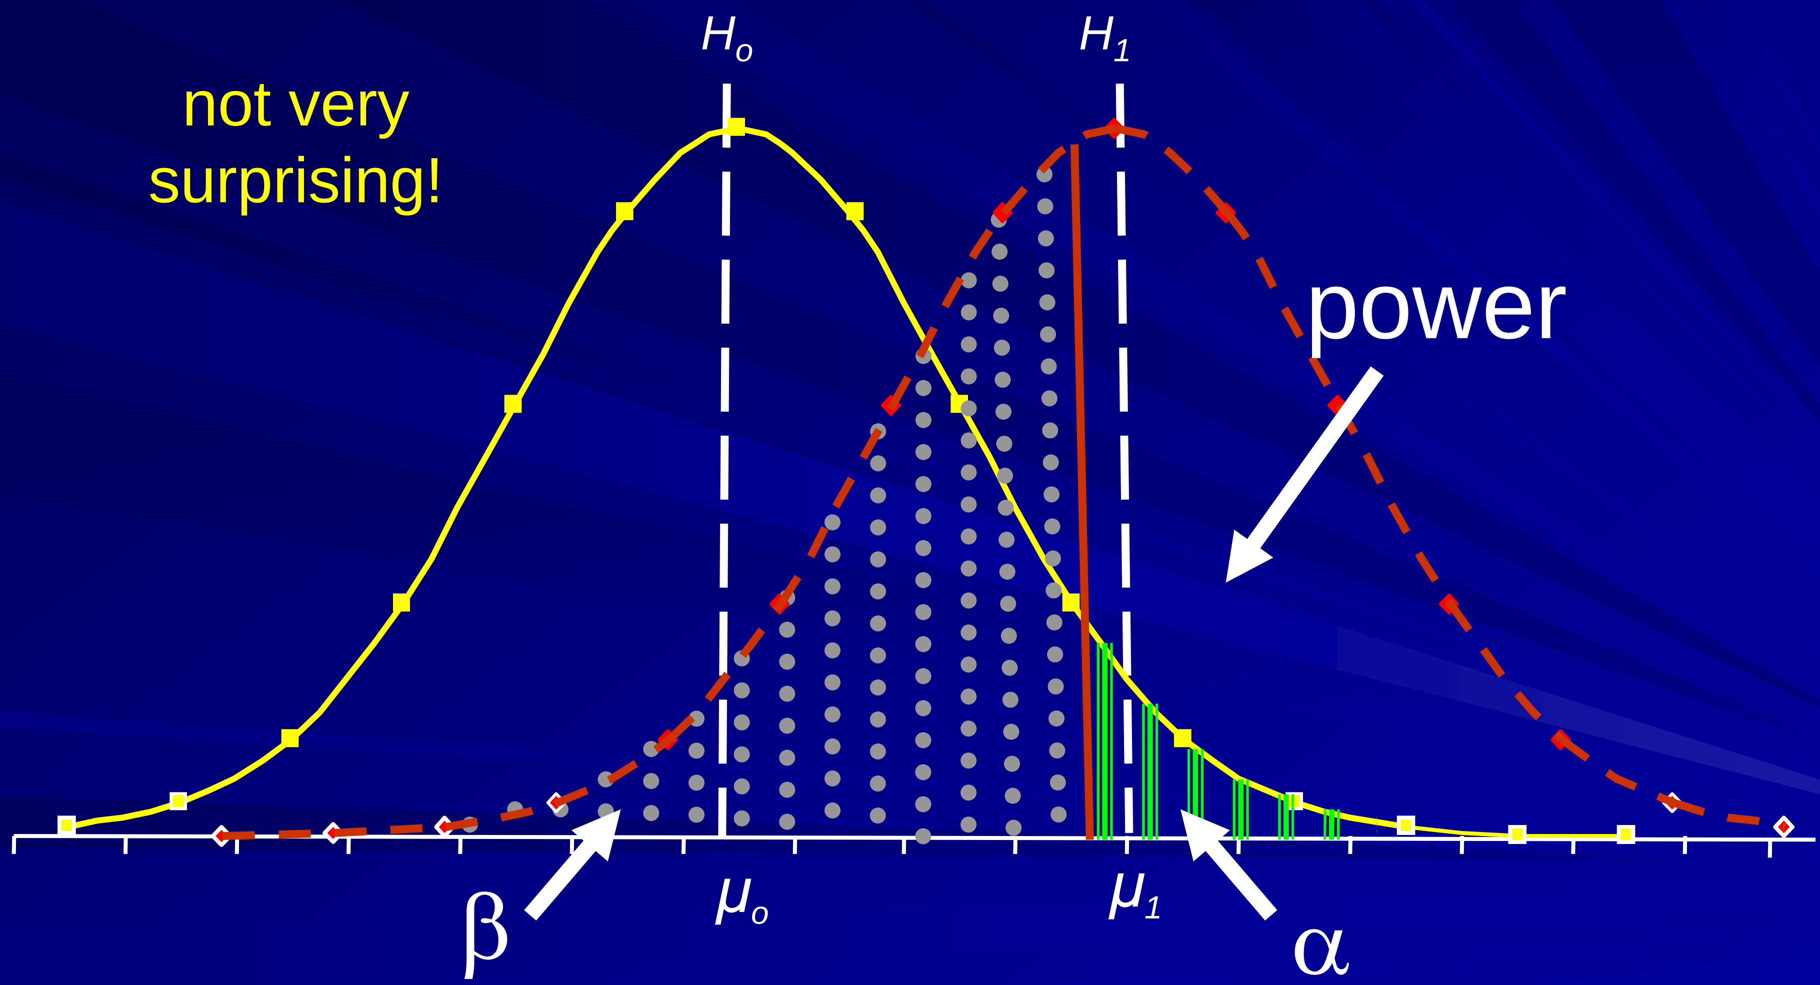

as a function of n and σ^2

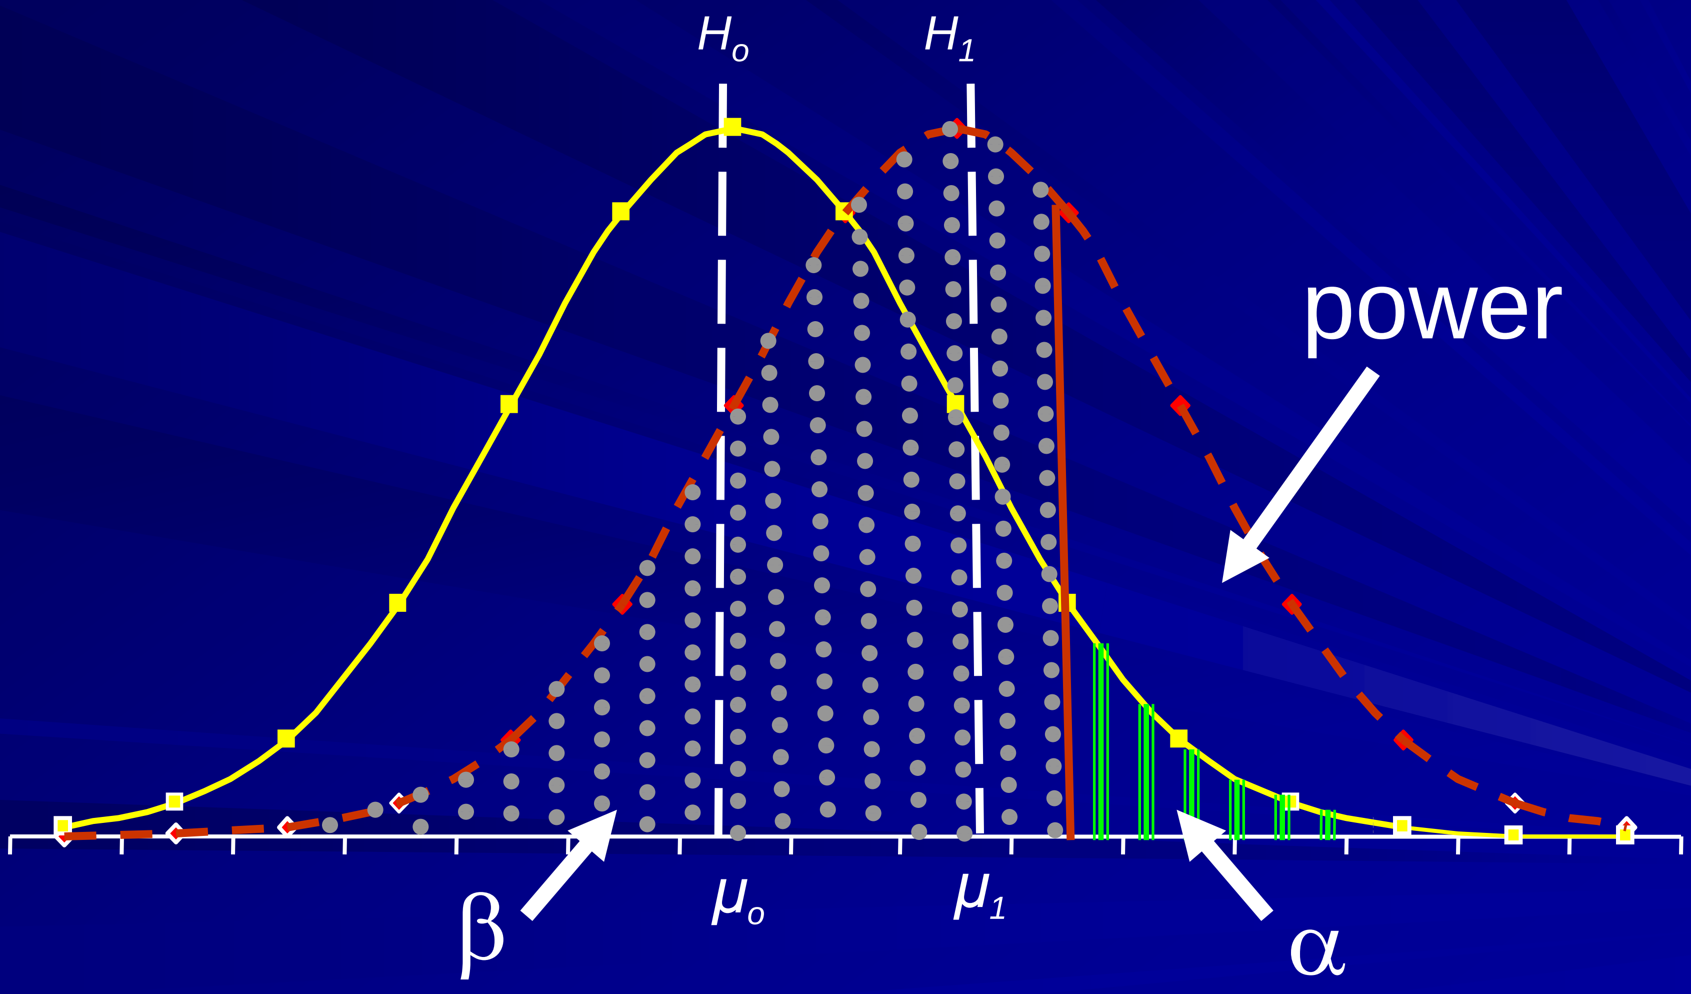

as a function of n and σ^2

an increase in sample size or decrease variance will decrease error variance

Recall: $\sigma^2_{means} = \sigma^2/n$

or $\sigma_{means} = \sigma^2/\sqrt{n}$

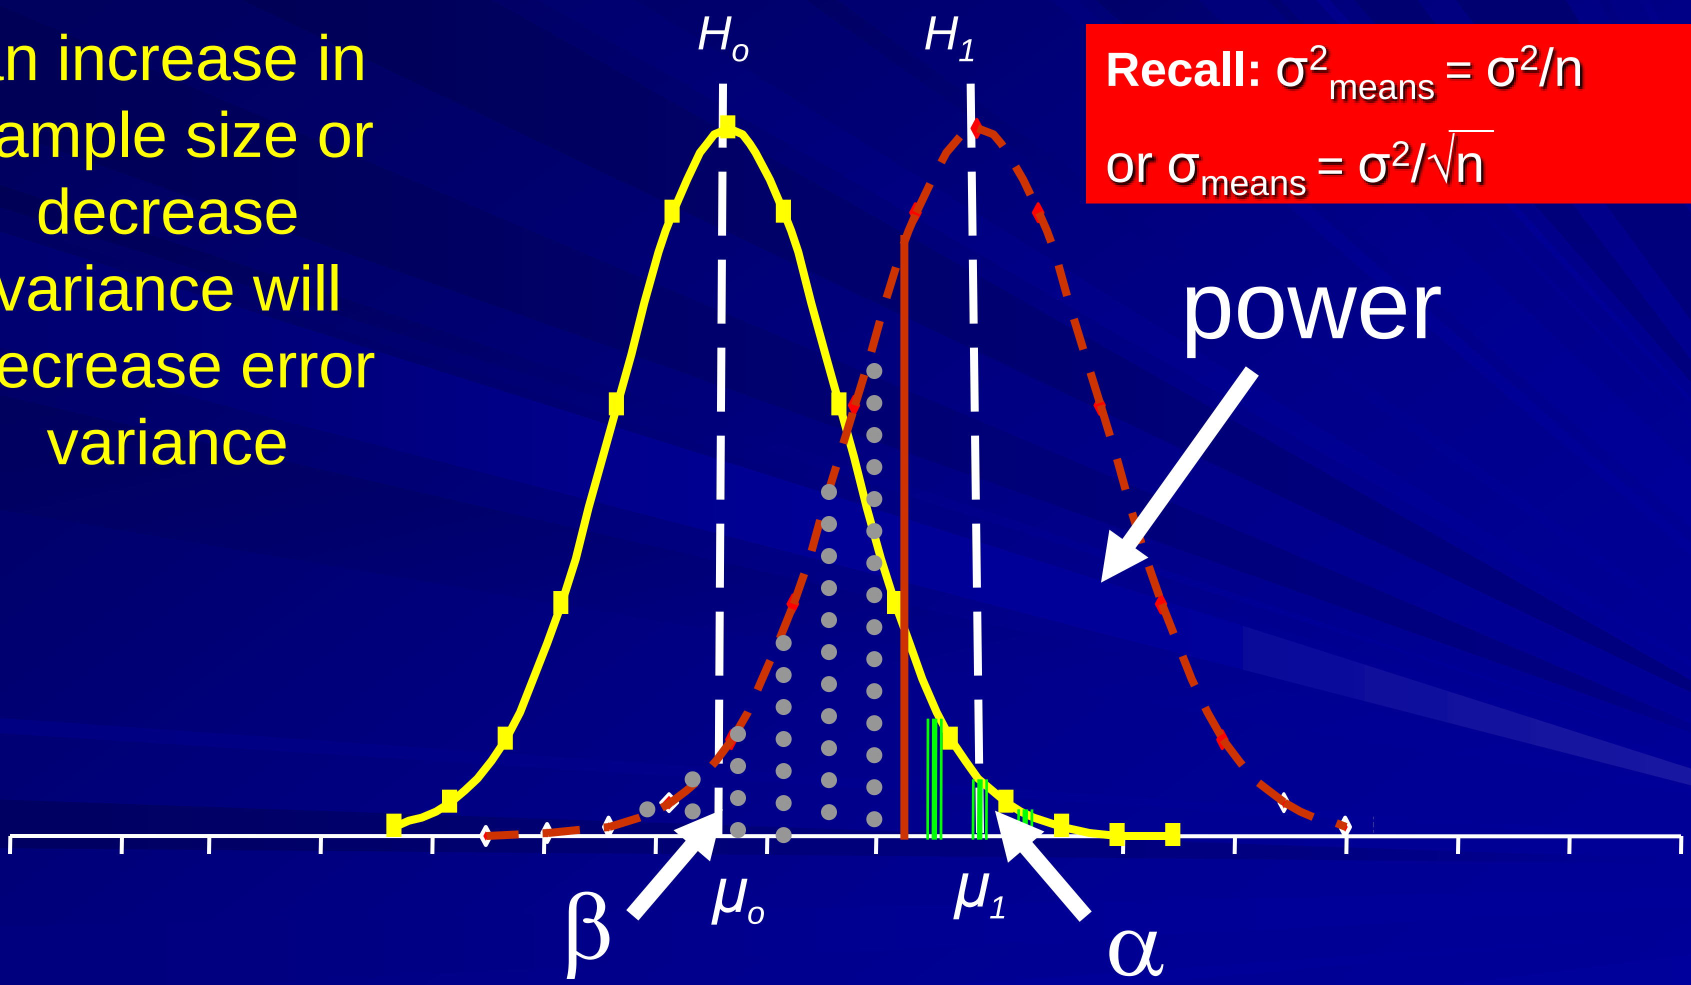

as a function of effect size (d)

- power closely related to effect size
- effect size is determined by some of these previous factors we looked at (μ_o - μ_1 and σ):
- $d = \dfrac{\mu_1 - \mu_0}{\sigma}$

i.e., $\dfrac{\text{how much the means differ}}{\text{variability in the sample}}$

- d indicates about how many standard deviations the means are apart, & thus the overlap of the two distributions:

effect size	d	percentage overlap
Small	0.20	85%
Medium	0.50	67%
Large	0.80	53%

put them all together…

- thus, we can calculate power if we know the effect size and the sample size
- for a t-test, this information can be combined using δ, a non-centrality parameter:

$$\delta = d\sqrt{n}$$

  - if we already have our data we know d and n
    - can then calculate power by computing δ and looking up tables *to **see what power we have*** (post hoc power)
  - if we don't have our data, we can specify δ and d
    - can then ***work out the n we need*** to get a certain level of power (a priori power)

Power & Effect

- Smaller effect sizes harder to find significant difference – need more powerful test
- "Magnifying glass"
- Big effect needs a less powerful

test (magnifying glass)

- Small effect needs a more

powerful test

…extend to 1-way anova…

- for anova the notation changes slightly:

$$\delta = d\sqrt{n} \quad \text{becomes} \quad \phi = \phi'\sqrt{n}$$

  - where n = number of observations in each treatment group
  - Φ′ = effect size (with 2 groups is same as d)
  - Φ′ calculation is conceptually similar to d:

$$\phi' = \sqrt{\frac{\sum \tau_j^2}{k\sigma_e^2}} = \sqrt{\frac{\sum (\mu_j - \mu)^2}{k\sigma_e^2}}$$

where μ = grand mean
μ_j = each group mean
k = number of groups
σ_e^2 = pooled variance (use MS_{error})

power for 1-way anova

- lets go back to our 'limericks' study…
  - suppose we didn't include the distraction factor – our data would look like this:

	mean creativity rating	n
0 pints of beer	63.75	16
2 pints of beer	64.69	16
4 pints of beer	46.56	16
$\boxed{MS_{error} = 125.21}$	$\mu = 58.33$	N=48

this is the error term we *would have* got running the study as a 1-way anova

power for effect of consumption

$$\phi' = \sqrt{\frac{\sum \tau_j^2}{k\sigma_e^2}} = \sqrt{\frac{\sum (\mu_j - \mu)^2}{k\sigma_e^2}}$$

where μ = grand mean = **58.33**

μ_j = each group mean = **63.75, 64.69, 46.56**

k = number of groups = **3**

σ_e^2 = pooled variance (use MS_{error}) = **125.21**

$$\phi' = \sqrt{\frac{(63.75 - 58.33)^2 + (64.69 - 58.33)^2 + (46.56 - 58.33)^2}{3 \times 125.21}}$$
$$= 0.56$$

$$\phi = \phi' \sqrt{n} = 0.56\sqrt{16} = 2.22$$

power for effect of consumption

- look up *Noncentral F (ncF)* tables in Howell, for $\Phi = 2.22$ with:
  - df_{treat} of 2 (k-1 = 3-1)
  - df_{error} of 45 (N-k = 48-3)
  - $\alpha = .05$
- this tells us the value of β, in this case 0.09
- as power = $1-\beta$, we have .91 power
  - that is, we have an 91% chance of detecting our effect (and correctly rejecting the null hypothesis), if the effect is real in the population
  - Cohen (1992) suggests >.80 is optimal

how about 2-way factorial anova?

formula is basically the same, only that now we have **3 effects** to estimate the power for:

$$\phi\alpha' = \sqrt{\frac{\sum \alpha_j^2}{a\sigma_e^2}} = \sqrt{\frac{\sum(\mu_j - \mu)^2}{a\sigma_e^2}}$$

$$\phi\alpha = \phi\alpha' \sqrt{nb}$$

$$\phi\beta' = \sqrt{\frac{\sum \beta_k^2}{b\sigma_e^2}} = \sqrt{\frac{\sum(\mu_k - \mu)^2}{b\sigma_e^2}}$$

$$\phi\beta = \phi\beta' \sqrt{na}$$

$$\phi\alpha\beta' = \sqrt{\frac{\sum \alpha\beta_{jk}^2}{ab\sigma_e^2}} = \sqrt{\frac{\sum(\mu - \mu_j - \mu_k + \mu_{jk})^2}{ab\sigma_e^2}}$$

$$\phi\alpha\beta = \phi\alpha\beta' \sqrt{n}$$

power for 2-way anova

- back to our 'limericks' study…
  - this time, we include the distraction factor:

	mean creativity rating		
	Distracted	Controls	(marginal mean)
0 pints of beer	66.88	60.63	***63.75***
2 pints of beer	66.88	62.50	***64.69***
4 pints of beer	35.63	57.50	***46.56***
(marginal mean)	***56.46***	***60.21***	

$\boxed{MS_{error} = 83.02}$ $\mu = 58.33$ N=48

can now use our error term from the 2-way anova

power for main effect of consumption

$$\phi\alpha' = \sqrt{\frac{\sum \alpha_j^2}{a\sigma_e^2}} = \sqrt{\frac{\sum(\mu_j - \mu)^2}{a\sigma_e^2}}$$

$$\phi\alpha' = \sqrt{\frac{(63.75 - 58.33)^2 + (64.69 - 58.33)^2 + (46.56 - 58.33)^2}{3 \times 83.02}}$$
$$= .915$$

$$\phi\alpha = \phi\alpha'\sqrt{nb}$$
$$= 0.915\sqrt{8x2}$$
$$= 3.66$$

(df_{treat} = 2 (3-1); df_{error} = 42 (48-6); α = .05)

this time β=0.00 (say 0.001) so power = .99

thus we have a 99% chance of detecting a significant effect if one exists in the pop.

(and correctly rejecting the null hypothesis)

what if we had <u>fewer</u> subjects?

- our main effect of consumption is so large that we would have enough power to detect it with much fewer subjects
- we could estimate what the power would be if we had done the study again with **3** people per cell (N=18):

$$\phi = \phi' \sqrt{nb} = 0.915 \times \sqrt{3 \times 2} = 2.24$$

df_{treat} of 2 (= 3-1)

  – df_{error} of 12 (18-6)
  – α = .05

- this time β = .14 so we have .86 power
  – that is, we have an 86% chance of detecting our effect
  (and correctly rejecting the null hypothesis, if effect exists in the pop)

power for main effect of distraction

$$\phi\beta' = \sqrt{\frac{\sum \beta_k^2}{b\sigma_e^2}} = \sqrt{\frac{\sum (\mu_k - \mu)^2}{b\sigma_e^2}}$$

$$\phi\beta' = \sqrt{\frac{(56.46 - 58.33)^2 + (60.21 - 58.33)^2}{2 \times 83.02}} = 0.206$$

$$\phi\beta = \phi\beta' \sqrt{na}$$
$$= 0.206\sqrt{8x3}$$
$$= 1.01$$

(df_{treat} = 1 (2-1); df_{error} = 42 (48-6); α = .05)

this time β = .71 so we have .29 power

that is, we have an 29% chance of detecting our effect

(and correctly rejecting the null hypothesis, if a real effect exists in the pop)

what if we had <u>more</u> subjects?

- our main effect of distraction is so small that we would need to have a lot more power to detect it
- we could estimate what the power would be if we had done the study again with **30** people per cell (N=360):

$$\phi\beta = \phi\beta'\sqrt{na}$$
$$= 0.206\sqrt{30x3}$$
$$= 1.95$$

$df_{treat} = 1(2\text{-}1)$; $df_{error} = 354\ (360\text{-}6)$; $\alpha = .05$

this time $\beta = .19$ so we have .81 power

that is, we have an 81% chance of detecting our effect

(and correctly rejecting the null hypothesis, if real diff exists in the pop)

What about higher-order ANOVA?

- Can generalize from this formula for each effect
- E.g. in 3-way ANOVA, look at power for each of the 7 effects (main effects of A, B, and C; AB, BC, and AC interactions; and ABC interaction)

$$\phi' = \sqrt{\frac{\sum \tau_j^2}{k\sigma_e^2}}$$

determining sample size

- if we can estimate our effect size and error we can plan studies in terms of how many participants we need to get .80 power
- often use previous studies to derive your estimates for MSE and effect size
- substituting different values of n into the formula and looking up the *ncF* tables is one (tedious) way of doing things…

G*POWER

- G*POWER is a FREE program that can make the calculations a lot easier

  http://www.psycho.uni-duesseldorf.de/abteilungen/aap/gpower3/

G*Power computes:

- power values for given sample sizes, effect sizes, and alpha levels (*post hoc power analyses*),
- sample sizes for given effect sizes, alpha levels, and power values (*a priori power analyses*)
- suitable for most major statistical methods
- V useful for hons thesis writing or report writing.

side issues…

- recall the logic of calculating estimates of **effect size** (i.e., criticisms of significance testing)
  - the tradition of significance testing is based upon an arbitrary rule leading to a yes/no decision
- power illustrates further some of the caveats with significance testing
  - with a high *N* you will have enough power to detect a very small effect
  - if you cannot keep error variance low a large effect may still be non-significant

side issues…

- on the other hand…
  - sometimes very small effects are important
  - by employing strategies to increase power you have a better chance at detecting these small effects

how can we maximize power?

- **power can be increased a number of different ways (some of these we have seen already…)**
  - increase sample size
  - increase α level (usually not wise!)
  - ***decrease error variance***
    - improving data collection methods:
      - procedural
      - psychometric
    - improving methods of analysis
      - i.e., use more *powerful* statistics

some methods we will look at…

methodological:

- blocking (next week)

statistical:

- analysis of covariance, ancova (next week)

- within-subjects anova (week 11 & 12)

Next week in class:

- Blocking and ANCOVA

Readings for this week:

- Howell
  - Chapter 16 (section 16.7)
  - Review chapter 9
- Field
  - Chapter 9

In the tutes:

- This week: SPSS tute
- Next week: Assignment 1 consult.